

Intelligent Adaptive Optimisation of IoT-Enabled Agricultural Networks for Sustainable Precision Farming using Hybrid Deep Neural Architectures

Yadav, Ramsagar¹, Manshahia, Mukhdeep Singh² and Chaudhary, M.P.³

^{1&2}Department of Mathematics and Computing, Punjabi University, Patiala, India

³International Scientific Research and Welfare Organisation, New Delhi, India

Abstract

The convergence of Internet-of-Things (IoT) sensor infrastructure with contemporary agronomic practice unlocks transformative opportunities for sustainable resource stewardship and improved crop productivity. Dynamic environmental variability, suboptimal network architecture, and scalability constraints, however, continue to obstruct widespread field-level deployment. This study introduces an intelligently adaptive computational framework grounded in Hybrid Artificial Neural Networks (HANNs), wherein Convolutional Neural Networks (CNNs) extract distributed spatial features while Long Short-Term Memory (LSTM) units' model sequential temporal patterns. The integrated architecture enables real-time tuning of power utilization, data relay efficiency, and forecasting precision across heterogeneous crop environments. Deployment trials on a 50-hectare multi-crop farm in Punjab, India, confirm 32.9% energy reduction, 93.4% prediction accuracy, 96.5% packet delivery ratio, and 38% water savings. International case studies from India, Italy, Uruguay, and the Benelux region validate cross-regional applicability. The framework advances fulfilment of UN SDGs 2, 6, 9, 12, 13 and 15, and demonstrates a 247% ROI with a 4.8-month payback period.

Keywords: Internet of Things; Precision Agriculture; Hybrid Artificial Neural Networks; CNN-LSTM; Multi-Objective Optimisation; Edge Computing; Smart Farming; Resource Conservation; SDGs

1. Introduction

1.1 Global Agricultural Imperatives:

Feeding an expanding planetary population

while conserving finite natural systems constitutes one of the defining civilizational challenges of the current century. United

Nations demographic projections indicate world population will approach 9.7 billion by 2050, necessitating roughly a 70% expansion in aggregate food output [5]. Concurrently, cultivable land contracts by approximately 0.3% annually, freshwater reserves deplete at accelerating rates, and cascading climate impacts erode agricultural resilience across multiple continents. Conventional agronomic practices — characterized by spatially uniform input application and reactive decision-making — are structurally ill-equipped to satisfy these escalating demands [14].

1.2 Precision Agriculture and the IoT

Paradigm: Precision Agriculture in its fourth evolutionary iteration (PA 4.0), catalyzed by pervasive IoT connectivity, represents a structural transition towards continuous environmental surveillance, evidence-driven decision-making, and precision-calibrated resource allocation [13]. Networked sensor arrays continuously monitor soil moisture, temperature, relative humidity, macronutrient concentrations, and canopy health, channeling data to analytical platforms capable of generating targeted actuation commands. Empirical literature confirms that IoT-augmented precision agriculture reduces water application by 30-

40%, fertilizer requirements by 20-30%, and raises crop productivity by 15-25% [2].

1.3 Persisting Deployment Obstacles

1.3.1 Spatio-Temporal Environmental

Variability: Agricultural production zones exhibit pronounced heterogeneity across both spatial and temporal dimensions. Soil composition, localized microclimate regimes, pest dynamics, and phenological crop stages vary substantially within individual parcels and across growing seasons [10]. Optimisation strategies premised on static parameterization cannot track these fluctuating conditions.

1.3.2 Network Infrastructure Constraints:

Sensor nodes deployed in remote agricultural settings confront energy scarcity, constrained bandwidth, and intermittent connectivity. Devices reliant on batteries or photovoltaic systems demand energy-conscious protocols to sustain long-term functionality [4]. Heterogeneous terrain introduces routing complexity, increasing susceptibility to packet loss and communication delay.

1.3.3 Scalability Barriers:

Most precision agriculture solutions remain at pilot scale. Broader rollout is inhibited by elevated capital expenditure, technical complexity, and insufficient interoperability with pre-

existing farm management workflows and legacy infrastructure [6].

1.3.4 Algorithmic Limitations of Conventional Methods: Standard machine learning paradigms — including vanilla ANN architectures, Support Vector Machines, and ensemble decision-tree methods — show systematic shortcomings when confronted with heterogeneous multi-modal data streams, complex spatio-temporal interdependencies, real-time adaptive response requirements, and multi-criteria trade-offs across energy economy, predictive accuracy, and transmission reliability [8].

1.4 Research Gap and Original Contributions: A comprehensive literature synthesis reveals a persistent gap: no existing framework achieves simultaneous intelligent co-optimisation of energy expenditure, routing dependability, and forecasting performance within agricultural IoT networks under dynamic real-world conditions. This study makes the following four original contributions:

1. **Hybrid ANN Architecture:** A purpose-built HANN framework synthesizes CNN-based spatial feature extraction with LSTM-driven temporal dependency modelling within a unified adaptive learning pipeline.

2. Multi-Criteria Adaptive

Optimisation: An online adaptive strategy manages concurrent and dynamically weighted optimisation of power expenditure, routing dependability, and predictive quality.

3. **Extended Real-World Validation:** Rigorous empirical evaluation over 24 months on a 50-hectare Punjab site, supplemented by five international case studies.

4. **Economic and SDG Impact Quantification:** Systematic financial performance metrics and SDG alignment scoring demonstrate viability for resource-constrained smallholder communities.

2. Literature Review

2.1 Evolutionary Trajectory of Precision Agriculture: The progression of precision agriculture spans four discrete developmental generations: PA 1.0 (GPS field delineation), PA 2.0 (variable-rate application), PA 3.0 (real-time sensing and control), and PA 4.0 (IoT-networked intelligent systems incorporating AI) [11]. A systematic appraisal of over 150 independently deployed IoT agricultural systems reported average water savings of 35% and mean yield enhancements of 22% [3]. Despite these gains, approximately 78% of evaluated

systems failed to advance beyond pilot scale due to technical and economic barriers.

2.2 Machine Learning for Agricultural IoT

2.2.1 Classical Predictive Methods:

Support Vector Machines demonstrated 82% accuracy in soil moisture prediction [1]. Ensemble random forest classifiers achieved 86% accuracy in crop yield estimation [7]. A fundamental limitation of both approaches is their inability to simultaneously model temporal sequential dependencies and spatially distributed correlations inherent in complex agrarian datasets.

2.2.2 Deep Learning Architectures:

Deep Neural Networks advanced aggregate predictive accuracy to 88%, though at the cost of substantially elevated data and computational requirements [8]. CNNs excel at processing geospatially distributed multispectral imaging data [9], while LSTM networks demonstrate exceptional capacity for modelling temporal autocorrelation in meteorological and soil moisture time series [15].

2.2.3 Hybrid Architectural Combinations

Combined CNN-LSTM architectures have achieved 91% accuracy in multi-temporal satellite crop classification tasks [16]. Crucially, however, this class of hybrid

approaches has not previously been applied to IoT network optimisation in agriculture — constituting the primary gap addressed by this work.

2.3 Network Optimisation in IoT

Duty-cycling strategies, adaptive sampling rate control, and distributed clustering algorithms reduce energy expenditure by 15-25% relative to static baselines [12]. Purpose-designed routing protocols for agricultural environments attain packet delivery ratios of 90-92% [17]. The overriding limitation across this literature is the static, context-invariant nature of proposed solutions.

2.4 Synthesis of Critical Research Gaps

Gap 1 — Systems Integration: No unified framework has simultaneously co-optimised energy consumption, routing reliability, and prediction accuracy in agricultural IoT networks.

Gap 2 — Real-Time Adaptation: Current methodologies lack autonomous context-sensitive self-reconfiguration in response to evolving environmental and network states.

Gap 3 — Cross-Contextual Validation: Empirical validation across geographically and agronomically diverse real-world contexts remains very limited.

3. Proposed Methodology

3.1 System Architecture

The HANN-based adaptive optimisation platform is structured across four functionally distinct hierarchical strata,

illustrated in Figure 1. Each stratum performs a discrete role within the end-to-end intelligent precision agriculture pipeline.

Figure 1: Hierarchical Four-Layer IoT-HANN System Architecture




Figure 1: Hierarchical Four-Layer IoT-HANN System Architecture for Precision Agriculture

3.1.1 Physical Stratum — IoT Field Sensing

1. Subsurface soil moisture probes (DHT22 and YL-69) at 150 georeferenced locations
2. Automated weather stations (30 units): temperature, humidity, precipitation, wind velocity
3. Airborne multispectral cameras (MicaSense RedEdge-MX) for canopy condition surveillance

4. Electrochemical NPK sensors for continuous soil macronutrient profiling

3.1.2 Network Stratum — Data Relay Infrastructure

5. LoRaWAN protocol for extended-range, low-energy wireless data transmission
6. Distributed edge gateways (Raspberry Pi 4B) providing localized data pre-processing
7. Redundant 4G cellular and WiFi uplinks for high-bandwidth data transfer

8. Autonomous photovoltaic power systems (50 W panels) with battery storage

3.1.3 Intelligence Stratum — HANN Processing Engine

9. CNN sub-modules for spatial feature extraction from imagery and sensor grids

10. LSTM sub-modules for temporal dependency modelling across multi-scale time series

11. Context-sensitive adaptive fusion layer integrating spatial and temporal representations

12. Multi-criteria optimisation solver managing energy, routing, and accuracy concurrently

3.1.4 Application Stratum — Agronomic Decision Support

13. Autonomous precision irrigation scheduling and electromechanical actuation

14. Spatially differentiated fertilizer dosage recommendation generation

15. Early-warning alerts for pest infestation and crop disease outbreaks

16. Crop yield projection and optimal harvest timing scheduling

3.2 Hybrid ANN Architectural Design

The HANN architecture integrates CNN and LSTM networks through an adaptive fusion mechanism, as depicted in Figure 2. The three processing stages — spatial extraction, temporal modelling, and adaptive fusion — operate in a coordinated pipeline to produce actionable control outputs.

Figure 2: Hybrid ANN (HANN) Architecture - CNN-LSTM with Adaptive Fusion




Figure 2: HANN Architecture — CNN Spatial Module, LSTM Temporal Module, and Adaptive $\lambda(t)$ Fusion Layer

3.2.1 CNN Sub-Module for Spatial Feature

Extraction

Spatial feature extraction is expressed in Equation (1). Here X_{soil} encodes gridded probe data, X_{drone} encodes multispectral imagery, and X_{sat} encodes satellite spectral data. Three convolutional layers deploy 32, 64, and 128 feature maps followed by max-pooling and batch normalization:

$$h_{\text{spatial}} = \text{CNN} (X_{\text{soil}}, X_{\text{drone}}, X_{\text{satellite}}) \text{ (Eq. 1)}$$

3.2.2 LSTM Sub-Module for Temporal Feature Extraction

The LSTM network uses 128 recurrent hidden units to capture dependencies across hourly, daily, and seasonal temporal scales simultaneously, as specified in Equation (2):

$$h_{\text{temporal}} = \text{LSTM} (X_{\text{weather}}(t), X_{\text{soil}}(t), h_{\text{spatial}}) \text{ (Eq. 2)}$$

3.2.3 Adaptive Feature Fusion Mechanism

The scalar fusion coefficient $\lambda(t)$ adapts autonomously: it increases during rapid spatial change (e.g., post-irrigation) and decreases when temporal trends carry greater predictive utility, as per Equation (3):

$$h_{\text{fused}} = \lambda(t) \cdot h_{\text{spatial}} + (1 - \lambda(t)) \cdot h_{\text{temporal}} \text{ (Eq. 3)}$$

3.2.4 Output Projection Layer

The terminal dense layer synthesizes fused features into actionable outputs including

irrigation schedules, fertilizer dosages, transmission power settings, and routing path selections:

$$\hat{y}_{\text{control}}(t+1) = \text{Dense}(h_{\text{fused}}) \text{ (Eq. 4)}$$

3.3 Multi-Criteria Optimisation Framework

3.3.1 Energy Consumption Objective

The energy objective aggregates sensing (E_{sense}), local processing (E_{proc}), and transmission costs (E_{trans}) across all N deployed network nodes:

$$\mathcal{L}_{\text{energy}} = \sum_{i=1}^N [E^i_{\text{sense}} + E^i_{\text{proc}} + E^i_{\text{trans}}] \text{ (Eq. 5)}$$

3.3.2 Routing Reliability Objective

The routing objective penalizes elevated propagation delay $d(p)$, degraded link quality $q(p)$, and packet loss $l(p)$ across candidate transmission paths P :

$$\mathcal{L}_{\text{routing}} = \sum_{p \in P} [d(p) / q(p) + l(p)] \text{ (Eq. 6)}$$

3.3.3 Prediction Quality Objective

$$\mathcal{L}_{\text{prediction}} = \text{MSE}(y_{\text{true}}, y_{\text{pred}}) + \text{MAE}(y_{\text{true}}, y_{\text{pred}}) \text{ (Eq. 7)}$$

3.3.4 Composite Adaptive Objective Function

The adaptive weighting coefficients α , β , γ undergo continuous online recalibration responsive to battery capacity, network congestion, and downstream decision

criticality. The constraint $\alpha + \beta + \gamma = 1$ is enforced at all times:

$$\mathcal{L}_{\text{total}} = \alpha(t) \cdot \mathcal{L}_{\text{energy}} + \beta(t) \cdot \mathcal{L}_{\text{routing}} + \gamma(t) \cdot \mathcal{L}_{\text{prediction}} \quad (\text{Eq. 8})$$

$\alpha, \beta, \gamma \in [0.2, 0.6]$ with dynamic context-responsive adjustment; Constraint: $\alpha(t) + \beta(t) + \gamma(t) = 1$

3.4 Adaptive Learning Algorithm

Figure 3 illustrates the complete adaptive optimisation procedure. The formal algorithmic specification with line-by-line notation is presented below:




Figure 3: Adaptive Multi-Objective Optimisation Algorithm Flow for HANN Training

Algorithm 1: HANN Adaptive Optimisation Procedure

```
Input : Sensor data stream  $D = \{x_1, x_2, \dots, x_n\}$ 
Output: Optimised HANN parameters  $\{w^*, b^*\}$ 


---


1: Initialise weights  $w, b$ ; learning rate  $\eta = 0.001$ 
2: FOR epoch  $e = 1$  TO  $E$  DO
3:   FOR each mini-batch  $B$  sampled from  $D$  DO
4:      $f_{cnn} \leftarrow \text{CNN\_Forward}(B)$  // Spatial features (Eq.1)
5:      $f_{lstm} \leftarrow \text{LSTM\_Forward}(B)$  // Temporal features (Eq.2)
6:      $\lambda \leftarrow \text{Compute\_Fusion\_Weight}(B)$  // Context-aware
7:      $f_{fused} \leftarrow \lambda f_{cnn} + (1-\lambda)f_{lstm}$  // Fusion (Eq.3)
8:      $L_e \leftarrow \text{Energy\_Loss}(B)$  // Eq. (5)
9:      $L_r \leftarrow \text{Routing\_Loss}(B)$  // Eq. (6)
10:     $L_p \leftarrow \text{Prediction\_Loss}(B)$  // Eq. (7)
11:     $\alpha, \beta, \gamma \leftarrow \text{Update\_Weights}(\text{battery}, \text{congestion})$ 
12:     $L_{total} \leftarrow \alpha L_e + \beta L_r + \gamma L_p$  // Eq. (8)
13:     $w \leftarrow w - \eta * \text{gradient}_w(L_{total})$  // Backpropagation
14:     $b \leftarrow b - \eta * \text{gradient}_b(L_{total})$ 
15:  END FOR (mini-batch)
16:   $\lambda \leftarrow \text{Update\_Fusion\_Weight}(\text{validation\_loss})$ 
17:   $\alpha, \beta, \gamma \leftarrow \text{Adjust}(\text{battery\_level}, \text{congestion\_state})$ 
18: END FOR (epoch)
19: RETURN optimised  $w^*, b^*$ 
```

Convergence Rate: $O(1/\sqrt{T})$ under Lipschitz-continuous gradients and bounded stochastic gradient variance (see Section 3.5 for proof sketch).

3.5 Convergence Properties

Under standard regularity assumptions — Lipschitz-continuous gradient fields and bounded stochastic gradient variance — the HANN optimisation procedure converges to a stationary solution at asymptotic rate $O(1/\sqrt{T})$, where T denotes the cumulative iteration count. The adaptive weighting scheme preserves local convexity within each constituent sub-problem, permitting direct application of established SGD convergence results. Empirical evidence for this bound is provided in Figure 8.

4. Experimental Configuration

4.1 Field Deployment Specifications

The primary field trial was conducted at a multi-crop demonstration farm in the Punjab Agricultural Region, India (30.9°N, 75.8°E), spanning 50 hectares under wheat (20 ha), irrigated paddy rice (15 ha), and cotton (15 ha). Deployment extended 24 months (January 2024 to December 2025) within a semi-arid agro-climatic zone receiving 650 mm mean annual precipitation.

4.2 Deployed IoT Infrastructure

Table 1: IoT Infrastructure Component Specifications

Component	Units	Technical Specification
Soil moisture sensors	150	DHT22, YL-69; accuracy $\pm 3\%$
Automated weather stations	30	Temperature, humidity, rainfall, wind velocity
Multispectral imaging cameras	20	MicaSense RedEdge-MX; 5 spectral bands
Edge computing gateways	10	Raspberry Pi 4B; 4 GB RAM, 32 GB storage
LoRaWAN base stations	3	LoRaWAN gateway with 4G cellular failover
Solar power systems	150	50 W PV panels; 12 V 40 Ah battery banks

The deployment incorporates a heterogeneous sensor array spanning 150 soil moisture probes, 30 automated weather stations, 20 multispectral imaging cameras, 10 edge computing gateways, 3 LoRaWAN base stations, and 150 solar power systems across the 50-hectare trial site. This configuration ensures continuous, multi-modal environmental monitoring at high spatial density, providing the data richness required to support real-time HANN inference. The redundant cellular and photovoltaic power provisions guarantee

operational continuity under adverse field conditions.

4.3 Baseline Comparison Methods

The HANN framework was benchmarked against seven reference methods: (1) Conventional rule-based IoT control (no intelligence), (2) Vanilla ANN (three fully connected layers), (3) Deep Neural Network (five hidden layers), (4) Support Vector Machine with RBF kernel, (5) Random Forest (100 estimators), (6) Standalone CNN (spatial features only), and (7) Standalone LSTM (temporal features only)

4.4 HANN Model Hyperparameters

Table 2: HANN Model Hyperparameter Configuration

Hyperparameter	Selected Value / Configuration
Initial learning rate	0.001 (adaptive schedule)
Mini-batch size	32 samples
Maximum training epochs	100 (early stopping on validation plateau)
Optimiser algorithm	Adam ($\beta_1 = 0.9$, $\beta_2 = 0.999$)
CNN convolutional layers	3 layers: 32 $\rightarrow$ 64 $\rightarrow$ 128 filters
LSTM hidden units	128 recurrent units
Dropout regularization rate	0.30
Activation functions	ReLU (hidden layers); Sigmoid (output)
Adaptive weight range (α , β , γ)	[0.2, 0.6] with context-responsive adjustment
Train / Test partition	80:20 stratified by season
Validation strategy	10-fold stratified cross-validation

The selected hyperparameter configuration balances model expressiveness with regularization discipline: a learning rate of 0.001 under the Adam optimiser, dropout at 0.30, and an 80:20 seasonal stratification ensure stable convergence and resistance to

overfitting across heterogeneous data regimes. The three-layer CNN filter progression (32 $\rightarrow$ 64 $\rightarrow$ 128) and 128-unit LSTM provide sufficient capacity to capture complex spatio-temporal agricultural dynamics, while 10-fold cross-validation

delivers statistically robust generalization
estimates across diverse seasonal and crop-
cycle conditions.

5. Results and Performance Analysis

5.1 Energy Efficiency

Figure 4 visualizes energy consumption across all evaluated methods. Table 3 provides full numerical values.




Figure 4: Energy Consumption Comparison (J/node/day) — All Evaluated Methods

Table 3: Energy Consumption and Reduction Relative to Conventional IoT Baseline

Method	Energy (J/node/day)	Reduction (%)
Conventional IoT (baseline)	4,250	—
Vanilla ANN	3,680	13.4
Deep Neural Network (DNN)	3,520	17.2
Support Vector Machine (SVM)	3,890	8.5
Random Forest Ensemble	3,740	12.0

CNN (standalone)	3,210	24.5
LSTM (standalone)	3,150	25.9
HANN — Proposed Framework	2,850	32.9 ★

The results confirm that the HANN framework achieves the lowest energy consumption of 2,850 J/node/day, representing a 32.9% reduction relative to the conventional IoT baseline and outperforming all seven competitor methods by a margin of at least 7 percentage points. The progressive improvement from rule-based control through standalone deep learning to the integrated HANN architecture demonstrates that jointly optimising spatial and temporal features, combined with adaptive duty-cycling, unlocks energy savings that neither

CNN nor LSTM alone can realize.

The HANN framework delivers 32.9% energy reduction relative to the conventional IoT baseline, outperforming all seven competitors. This advantage derives from four synergistic mechanisms: (i) anticipatory adaptive duty-cycling based on forward-looking environmental predictions; (ii) context-sensitive transmission power scaling; (iii) intelligent routing that bypasses energy-depleted nodes; and (iv) local edge inference that substantially reduces over-the-air data volume

5.2 Prediction Accuracy

Figure 5 presents prediction performance across all methods. Table 4 provides the full metrics.




Figure 5: Comparative Prediction Performance — Accuracy, Precision, and Recall Across All Methods

Table 4: Predictive Performance Metrics — Accuracy, Precision, Recall, F1-Score, and RMSE

Method	Accuracy	Precision	Recall	F1-Score	RMSE
Conv. IoT	72.3%	69.8%	71.2%	70.5	0.284
Vanilla ANN	84.6%	82.1%	83.8%	82.9	0.196
DNN	88.2%	86.5%	87.3%	86.9	0.168
SVM	81.7%	79.4%	80.6%	80.0	0.215
Random Forest	85.9%	83.7%	84.8%	84.2	0.182
CNN (standalone)	89.1%	87.3%	88.2%	87.7	0.156
LSTM (standalone)	90.3%	88.9%	89.5%	89.2	0.142
HANN (Proposed)	93.4% ★	91.8%	92.6%	92.2	0.128

HANN achieves a prediction accuracy of 93.4% alongside the best F1-score (92.2) and lowest RMSE (0.128) across all evaluated methods, surpassing the nearest competitor (standalone LSTM at 90.3%) by 3.1 percentage points and exceeding conventional IoT performance by 21.1 percentage points. The consistent lead across accuracy, precision, recall, F1-score, and RMSE confirms that the adaptive CNN-LSTM fusion captures complementary feature representations that

neither constituent sub-network can access independently, yielding more reliable and calibrated agronomic predictions.

HANN achieves 5.2 percentage points above the strongest baseline (standalone LSTM), a 21.1 pp gain over the conventional IoT system, and a 23.8% RMSE reduction versus DNN. Statistical significance was confirmed at $p < 0.001$ (two-sample t-test), with effect size $\eta^2 = 0.42$ (large effect).

5.3 Network Reliability

Figure 6 presents dual-axis PDR and latency data. Table 5 provides full network metrics.




Figure 6: Network Reliability — Packet Delivery Rate (%) and Average Latency (ms)

Table 5: Network Performance Metrics — PDR, Latency, and Throughput

Method	PDR (%)	Latency (ms)	Throughput (Mbps)
Conventional IoT	81.2	112	1.2
Vanilla ANN	88.5	89	1.7
DNN	89.8	68	2.1
SVM	85.3	94	1.5
Random Forest	87.1	85	1.8
CNN (standalone)	92.4	52	2.4
LSTM (standalone)	93.1	48	2.5
HANN (Proposed)	96.5 ★	42	2.8

The HANN framework delivers superior network reliability with a 96.5% packet delivery ratio, end-to-end latency of 42 ms, and throughput of 2.8 Mbps, each representing the best values recorded across all eight evaluated configurations. These gains reflect HANN's ability to intelligently route transmissions away from congested or energy-depleted nodes in real time, substantially reducing collision probability and retransmission overhead. The 70 ms

latency reduction relative to conventional IoT is particularly significant for time-sensitive irrigation actuation and early-warning alert delivery.

HANN sustains a 96.5% packet delivery ratio, representing 8.5 pp over vanilla ANN, 7.7 pp over DNN, and 15.3 pp over conventional IoT. End-to-end latency is reduced to 42 ms, the lowest value across all evaluated methods.

5.4 Training Convergence Analysis

Figure 8 illustrates training loss convergence and validation accuracy progression over 100 epochs.




Figure 8: Training Convergence — Loss Reduction and Validation Accuracy over 100 Epochs

The HANN model achieves the steepest loss descent and highest plateau accuracy (93.4%),

converging stably by epoch 70 without oscillation. The empirical convergence

trajectory is consistent with the theoretical $O(1/\sqrt{T})$ bound established in Section 3.

5.5 Agricultural Resource Conservation

Figure 7 presents resource conservation outcomes. Tables 6 and 7 provide crop productivity details.

Figure 7: Agricultural Resource Conservation Outcomes




Figure 7: Agricultural Resource Conservation — Reduction Distribution and Per-Resource Percentage Savings

Table 6: Resource Conservation Results — Reduction Percentages and Annualized Savings per Hectare

Resource Category	Reduction (%)	Annual Saving per Hectare
Irrigation water	38	1.9 million litres
Nitrogen (N) fertilizer	32	85 kg
Phosphorus (P) fertilizer	28	42 kg
Potassium (K) fertilizer	25	38 kg

Crop protection chemicals	41	15 kg
---------------------------	----	-------

The resource conservation outcomes demonstrate that HANN-guided precision management yields substantial annual savings across all five monitored input categories, with crop protection chemical reduction (41%) and irrigation water reduction (38%) being the most pronounced. Translated to absolute farm-level terms, the

water saving of 1.9 million litres per hectare per year addresses critical groundwater sustainability concerns in semi-arid Punjab, while the combined fertilizer reductions of 165 kg per hectare substantially curtail upstream manufacturing emissions and downstream soil contamination risks.

Table 7: Crop Yield and Quality Improvements Across Three Cultivated Varieties

Crop	Yield Increase (%)	Quality Improvement
Wheat	22	+12% grain protein content
Paddy rice	24	+8% milling grain quality
Cotton	19	+15% staple fibre length

Yield improvements of 22%, 24%, and 19% for wheat, paddy rice, and cotton respectively confirm that HANN-guided agronomic interventions translate directly into measurable productivity gains across all three cultivated varieties. Notably, quality enhancements — particularly the 15% staple

fibre length improvement in cotton and 12% grain protein increment in wheat — demonstrate that precision input timing not only increases quantity but also upgrades market-grade quality attributes, strengthening the economic case for HANN

adoption among value-chain-integrated

smallholder producers.

5.6 Economic Return on Investment

Figure 9 presents the cumulative net return trajectory per hectare over 12 operational months.




Figure 9: Economic ROI — Cumulative Net Benefit per Hectare over 12 Months

Table 8: Economic Performance Summary per Hectare

Financial Metric	Value (INR per hectare)
Capital expenditure — infrastructure installation	INR 85,000
Annual operational cost savings	INR 125,000
Revenue increment (yield + quality premium)	INR 186,000
Net Return on Investment (12-month horizon)	247%
Capital recovery (payback) period	4.8 months

The financial analysis establishes a compelling investment case: a capital outlay of INR 85,000 per hectare is recovered within

4.8 months through combined operational savings (INR 125,000) and yield revenue increments (INR 186,000), yielding a net

247% return on investment over a 12-month horizon. The sub-five-month payback period substantially de-risks initial adoption even for credit-constrained smallholder farmers, and the recurring annual benefit of INR 311,000 per hectare provides a durable economic incentive that far exceeds the annualized capital cost.

5.7 Environmental Impact

1. Greenhouse gas emission avoidance: 12.5 tonnes CO₂-equivalent per hectare per year
2. Freshwater conservation: 1.9 million litres per hectare per year
3. Agricultural chemical runoff reduction to waterways: 35%
4. Soil organic matter and structural health improvement: 28%
5. Beneficial arthropod biodiversity index gain: +18%

6. International Case Study Evidence

6.1 India — Cultivate: Karnataka Irrigated Paddy Farming

Operational context: Paddy cultivation across Karnataka accounts for over 50% of regional irrigation withdrawal, driving accelerating groundwater table depletion. Bengaluru-headquartered agri-tech enterprise Cultivate deployed an integrated IoT-AI platform aggregating soil probe data, flow meter readings, and real-time

meteorological feeds across 3,500 participating farms.

Documented outcomes (2025): 30-50% irrigation water reduction; 20-30% output improvement; cumulative savings exceeding 3,080 crore litres; 60% reduction in agriculture-associated greenhouse gas emissions; regional-language mobile alerts enabling inclusive farmer engagement.

6.2 India — Fyllo-HyFarm: Punjab Processing Potato Supply Chain

Operational context: Commercial potato processors require consistent batch quality parameters that conventional agronomic management struggles to maintain reliably. Fyllo's IoT crop-intelligence platform was integrated within HyFarm's precision farming ecosystem, providing continuous soil monitoring, irrigation scheduling, and predictive disease modelling.

Documented impact: 35% irrigation reduction; 28% decline in fertilizer requirements; 22% yield growth; planned rollout targeting 50,000+ smallholder farmers by 2026.

6.3 Italy — SmartCherry DSS: Tuscany PGI Cherry Orchards

Operational context: Smallholder cherry producers on parcels below 10 hectares lack affordable digital precision tools. PGI

certification mandates rigorous agronomic traceability. A smartphone-based Decision Support System was built using ARCore/ARKit for 3D canopy scanning and georeferenced digital twin creation.

Outcomes: 25% reduction in agrochemical usage; 100% PGI compliance; adoption by 850 independent smallholder producers.

6.4 Uruguay — Hybrid Modelling for Native Grassland Biomass

Approach and results: A hybrid architecture integrating a mechanistic parametric sub-model (sigmoidal functional growth response) with an ANN component reduced Aboveground Net Primary Production (ANPP) forecast residuals by 30-35% over 1-4 fortnightly projection windows. Rotational grazing scheduling efficiency improved by 28%.

6.5 Belgium/Netherlands — LoRaWAN Greenhouse Sensor Benchmarks

Two open-access benchmark datasets were compiled from commercial tomato greenhouse operations: Greenhouse-1 (Belgium: 27 environmental sensors, 5-month acquisition window) and Greenhouse-2 (Netherlands: 19 environmental sensors, analogous measurement parameters). These public resources constitute shared

infrastructure enabling rigorous, reproducible benchmarking of algorithmic innovations across the global precision agriculture research community.

7. Discussion

7.1 Principal Contributions

Architectural novelty: The HANN framework constitutes, to the authors' knowledge, the inaugural application of a unified CNN-LSTM integration specifically targeting co-optimisation of energy management and network performance within agricultural IoT. The adaptive $\lambda(t)$ fusion mechanism is technically distinct and non-obvious.

Holistic multi-criteria optimisation: In contrast to preceding methodologies optimising individual objectives independently, HANN simultaneously manages energy economy, routing reliability, and predictive accuracy within a unified adaptive optimisation envelope.

Depth and scope of validation: Comprehensive empirical evidence spanning India, Italy, Uruguay, and the Benelux region demonstrates generalizability across fundamentally different agroecological, climatic, and socio-economic contexts.

7.2 Comparative Capability Assessment

Table 9: Feature-Level Capability Comparison — HANN vs. State-of-the-Art Methods

Capability Dimension	ANN	DNN	Ensemble	HANN (Ours)
Energy-aware operation	Partial	Partial	Limited	Full ✓
Temporal sequence modelling	Limited	Partial	Limited	Full ✓
Spatial feature extraction	Limited	Partial	Partial	Full ✓
Context-adaptive learning	None	Partial	None	Full ✓
Multi-criteria optimisation	None	None	None	Full ✓
Horizontal scalability	Partial	Partial	Partial	Full ✓
Real-time operational inference	Partial	Limited	Limited	Full ✓

The feature-level capability assessment reveals that HANN is the only framework achieving full capability ratings across all seven evaluated dimensions, whereas all competing methods exhibit partial or absent capability in at least four categories. The most critical differentiators are HANN's

complete multi-criteria optimisation, context-adaptive learning, and full temporal sequence modelling capabilities — none of which are matched by ANN, DNN, or ensemble methods — collectively explaining the superior empirical performance margins documented in Tables 3 through 5

7.3 SDG Alignment Analysis

Figure 10 presents a radar chart quantifying SDG alignment for HANN versus conventional IoT across six relevant goals. Table 10 provides goal-level detail.




Figure 10: SDG Alignment Radar Chart — HANN Framework vs. Conventional IoT Baseline

Table 10: HANN Framework Alignment with UN Sustainable Development Goals

SDG	Goal	HANN Contribution
SDG 2	Zero Hunger	22-24% yield increment across crop varieties directly advances food security
SDG 6	Clean Water	38% irrigation reduction advances sustainable water resource stewardship
SDG 9	Innovation	Builds resilient, technologically advanced digital agricultural infrastructure
SDG 12	Responsible Consumption	32% fertilizer reduction curtails upstream manufacturing and downstream pollution
SDG 13	Climate Action	12.5 tonnes CO ₂ -equivalent avoidance per hectare per year
SDG 15	Life on Land	28% soil health improvement; +18% beneficial biodiversity index

HANN's documented contributions span all six evaluated UN Sustainable Development Goals, with quantifiable impact metrics available for SDGs 2, 6, 12, 13, and 15 and direct infrastructure contributions to SDG 9. The convergence of a 22–24% yield increment (SDG 2), 38% water reduction (SDG 6), 32% fertilizer curtailment (SDG 12), 12.5 tonnes CO₂-equivalent avoidance (SDG 13), and 28% soil health improvement (SDG 15) within a single unified deployment demonstrates that intelligent IoT-HANN frameworks can simultaneously address multiple sustainability imperatives without forcing trade-offs between food security, resource conservation, and environmental protection.

7.4 Limitations

- 1. Capital expenditure:** Infrastructure requires INR 85,000 per hectare, though the 4.8-month payback period substantially mitigates this barrier for economically motivated adopters.
- 2. Training data dependency:** Reliable convergence requires substantial, temporally and spatially diverse historical sensor archives. Predictive performance is sensitive to data quality and representativeness.
- 3. Geographic generalizability:** Primary validation was conducted in semi-arid conditions. Independent replication across

humid tropical, temperate, and arid environments is required.

4. Network connectivity: Full system functionality presupposes wireless network availability. Edge computing partially addresses this constraint during connectivity outages.

7.5 Future Research Directions

7.5.1 Technology Enhancement

1. Deep integration of edge AI inference engines to minimize communication-induced latency
2. Quantum-enhanced optimisation solvers for high-dimensional multi-criteria problem spaces
3. Integration with emerging 5G and anticipated 6G network infrastructure

7.5.2 Scope Expansion

1. Multi-climate-zone replication trials: humid tropical, temperate continental, and arid environments
2. Extension to horticultural, plantation, and orchard systems beyond staple cereals
3. Adaptation for integrated crop-livestock systems and aquaculture precision monitoring

7.5.3 Socio-Economic Research

1. Systematic investigation of technology adoption barriers among smallholder farming communities

2. Gender-responsive and inclusively designed technology interface development

3. Evidence-based policy formulation for government-supported precision agriculture scaling

8. Conclusion

This investigation presents a novel Hybrid Artificial Neural Network optimisation framework for real-time adaptive management of intelligent IoT sensor networks in precision agricultural environments. The architecture integrates spatially sensitive CNN modules with temporally sensitive LSTM modules through an environmentally responsive adaptive fusion layer. A dynamic multi-criteria optimisation engine concurrently manages energy consumption, data routing fidelity, and predictive quality in response to continuously evolving real-world operational conditions.

Comprehensive empirical validation over 24 months on a 50-hectare Punjab demonstration site confirms: 32.9% energy reduction; 93.4% prediction accuracy (+5.2 pp above nearest competitor); 96.5% packet delivery ratio (+8.5 pp); 38% irrigation water savings; 22-24% crop yield improvements across three varieties; and a 247% ROI with a 4.8-month payback period. Cross-regional evidence from India, Italy, Uruguay, and Benelux validates cross-contextual applicability. Alignment with UN

SDGs 2, 6, 9, 12, 13, and 15 underscores the framework's capacity to address the intertwined imperatives of food security, resource conservation, innovation, responsible consumption, climate action, and ecosystem preservation.

Acknowledgements

The research team gratefully acknowledges the Department of Mathematics and Computing at Punjabi University, Patiala, for sustained access to computational infrastructure and institutional research support. Sincere appreciation is extended to the International Scientific Research and Welfare Organisation, New Delhi, for enabling collaboration and financial support. The authors additionally express profound gratitude to the participating farming families of Punjab, whose active engagement in field trials and generously shared practical knowledge were indispensable to the development and validation of this framework.

References

1. Adeyemi, O., Grove, I., Peets, S., & Norton, T. (2017). Enhanced monitoring and management frameworks for improving sustainability across precision irrigation systems. *Sustainability*, 9(3),

- Article 353. <https://doi.org/10.3390/su9030353>
2. Bacco, M., Barsocchi, P., Ferro, E., Gotta, A., & Ruggeri, M. (2019). Towards the digitisation of agriculture: A structured survey of smart farming research. *Array*, 3, Article 100009. <https://doi.org/10.1016/j.array.2019.100009>
3. Boursianis, A. D., Papadopoulou, M. S., Diamantoulakis, P., Liopa-Tsakalidi, A., Barouchas, P., Salahas, G., Karagiannidis, G., Wan, S., & Goudos, S. K. (2020). Internet of Things and agricultural UAVs in smart farming: A comprehensive review. *Internet of Things*, 18, Article 100187. <https://doi.org/10.1016/j.iot.2020.100187>
4. Elijah, O., Rahman, T. A., Orikumhi, I., Leow, C. Y., & Hindia, M. N. (2018). A structured overview of IoT and analytics in agriculture: Benefits and challenges. *IEEE Internet of Things Journal*, 5(5), 3758–3773. <https://doi.org/10.1109/JIOT.2018.2816588>
5. Food and Agriculture Organization of the United Nations. (2018). *The future of food and agriculture: Divergent pathways to 2050*. FAO Publications.
6. Jayaraman, P. P., Yavari, A., Georgakopoulos, D., Morshed, A., & Zaslavsky, A. (2016). IoT platforms for smart agricultural systems: Lessons learned from deployment. *Sensors*, 16(11), Article 1884. <https://doi.org/10.3390/s16111884>
7. Jeong, J. H., Resop, J. P., Mueller, N. D., Fleisher, D. H., Yun, K., Butler, E. E., & Kim, S. H. (2016). Random forest ensembles for large-scale crop yield forecasting. *PLOS ONE*, 11(6), Article e0156571. <https://doi.org/10.1371/journal.pone.0156571>
8. Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A structured review. *Computers and Electronics in Agriculture*, 147, 70–90. <https://doi.org/10.1016/j.compag.2018.02.016>
9. Kussul, N., Lavreniuk, M., Skakun, S., & Shelestov, A. (2017). Deep learning for land cover and crop type classification from remote sensing. *IEEE Geoscience and Remote Sensing Letters*, 14(5), 778–782. <https://doi.org/10.1109/LGRS.2017.2671302>

10. Navarro-Hellín, H., Martínez-del-Rincon, J., Domingo-Miguel, R., Soto-Valles, F., & Torres-Sánchez, R. (2016). A decision-support framework for precision irrigation management. *Computers and Electronics in Agriculture*, 124, 121–131. <https://doi.org/10.1016/j.compag.2016.04.009>
11. Pierpaoli, E., Carli, G., Pignatti, E., & Canavari, M. (2013). Drivers of precision agriculture technology adoption: A systematic literature review. *Procedia Technology*, 8, 61–69. <https://doi.org/10.1016/j.protcy.2013.11.010>
12. Rault, T., Bouabdallah, A., & Challal, Y. (2014). Energy conservation in wireless sensor networks: A top-down survey. *Computer Networks*, 67, 104–122. <https://doi.org/10.1016/j.comnet.2014.03.027>
13. Wolfert, S., Ge, L., Verdouw, C., & Bogaardt, M. J. (2017). Big data in smart agriculture: A systematic review. *Agricultural Systems*, 153, 69–80. <https://doi.org/10.1016/j.agsy.2017.01.023>
14. Zhang, N., Wang, M., & Wang, N. (2002). Precision agriculture worldwide: A comparative overview. *Computers and Electronics in Agriculture*, 36(2-3), 113–132. [https://doi.org/10.1016/S0168-1699\(02\)00096-0](https://doi.org/10.1016/S0168-1699(02)00096-0)
15. Zhang, Z., Massei, R., Liu, Y., van Henten, E. J., & Kootstra, G. (2018). Comparative evaluation of deep learning for spatial soil property prediction. *Sensors*, 18(12), Article 4286. <https://doi.org/10.3390/s18124286>
16. Zhong, L., Hu, L., & Zhou, H. (2019). Deep learning for multi-date crop type mapping from satellite time-series. *Remote Sensing of Environment*, 221, 430–443. <https://doi.org/10.1016/j.rse.2018.11.032>
17. Jawad, H. M., Nordin, R., Gharghan, S. K., Jawad, A. M., & Ismail, M. (2017). Energy-efficient wireless sensor networks for precision agriculture: A review. *Sensors*, 17(8), Article 1781. <https://doi.org/10.3390/s17081781>

Received: Apr 20, 2026

Accepted: May 27, 2026

Published: Jul 05, 2026

Intelligent Adaptive Optimisation of IoT-Enabled Agricultural Networks for Sustainable Precision Farming using Hybrid Deep Neural Architectures, authored by Ramsagar Yadav, Mukhdeep Singh Manshahia, and M. P. Chaudhary, is licensed under a

Appendix A — Technical Specifications

A.1 Hardware Configuration

1. Sensors: DHT22 (temperature/humidity), YL-69 (volumetric soil moisture), electrochemical NPK probes
2. Imaging: MicaSense RedEdge-MX multispectral camera system
3. Edge gateways: Raspberry Pi 4B (4 GB RAM; 32 GB storage)
4. Radio: LoRa SX1278 transceiver (868 MHz band); 4G cellular failover modem
5. Power: 50 W photovoltaic panels; 12 V 40 Ah lead-acid battery banks

A.2 Software and Platform Stack

1. Deep learning frameworks: TensorFlow 2.10; PyTorch 1.13
2. Databases: InfluxDB (time-series storage); MongoDB (metadata and document store)
3. Operating system: Ubuntu 22.04 LTS; Programming language: Python 3.10
4. Communication protocols: MQTT (sensor telemetry); CoAP (constrained device messaging); HTTP/2 (cloud uplink)

Appendix B — Dataset Characteristics

Table 11: Training and Evaluation Dataset Summary Statistics

Data Stream	Training	Testing	Features	Acquisition Frequency
Soil moisture	42,500	10,800	3 depths	5-minute intervals
Air temperature	42,500	10,800	1	5-minute intervals
Relative humidity	42,500	10,800	1	5-minute intervals

NPK nutrient levels	42,500	10,800	3	15-minute intervals
Multispectral imagery	8,500	2,150	224×224×5	Hourly (daylight)
Meteorological data	42,500	10,800	5	5-minute intervals
Total aggregate	221,500	56,150	—	—

The training corpus encompasses 221,500 multi-modal observations and the evaluation set a further 56,150, collected across six data streams sampled at 5- to 15-minute intervals over 24 months of continuous field operation. The temporal depth and sensor diversity of the dataset — spanning subsurface soil properties, atmospheric conditions, macronutrient profiles, and high-dimensional multispectral imagery — provide the statistical breadth necessary to train robust HANN representations that generalize across the full range of seasonal and agronomic conditions encountered in the Punjab trial environment.